Gigamon®

2026 HYBRID CLOUD SECURITY SURVEY

Reality Check: Exposing the AI Security Illusion

3 Introduction: A False Sense of Security

4 Key Findings

5 Confidence Without Control

7 The Visibility Gap Behind the Illusion

9 When Hybrid Cloud Risk is Easy to Misread

10 The AI Security Illusion Reaches the Boardroom

12 CISO Perspective: From Confidence to Proof

13 From Illusion to Evidence

14 Methodology

15 References

A False Sense of Security

Over the past year, AI has quickly moved from emerging priority to operating reality. It is no longer confined to experimentation or isolated use cases. It is being embedded within applications, workflows, and infrastructure at a pace that is reshaping how organizations operate and how data moves across hybrid cloud environments.

This acceleration is not incremental. It is structural.

Global investment in AI has risen sharply in recent years, with worldwide AI spending reaching a staggering[1] **$1.5 trillion** in 2025, and organizations are racing to operationalize and quantify its value across every business function. As AI adoption expands, so does the volume, velocity, and complexity of data in motion. Applications are more distributed, interactions are more dynamic, and dependencies are harder to trace. The modern enterprise is becoming more intelligent, but concerningly more opaque.

To keep pace, organizations are investing heavily in security with more tools, more policies, and more oversight. Governance frameworks are evolving, and executive attention on cyber risk is higher than ever. On the surface, it appears security is keeping up with the speed of change.

But the data reveals a different and disturbing reality.

Insights from more than 1,000 Security and IT leaders across six global regions, including over 300 CISOs, reveal a consistent and concerning pattern. **Nearly two-thirds** of organizations experienced a breach in the past year, an **18 percent** increase year over year, and nearly **one-third** (32 percent) reported multiple incidents. Yet **1 in 4** (27 percent) could not determine root cause, and only **30 percent** say they have the tools required to respond effectively.

At the same time, AI is now deeply embedded in the threat landscape. Today, **83 percent** of organizations report AI involvement in their security incidents, spanning external attacks, internal exposure, and the growing targeting of AI systems themselves.

Yet still, organizations express high confidence in their AI security maturity.

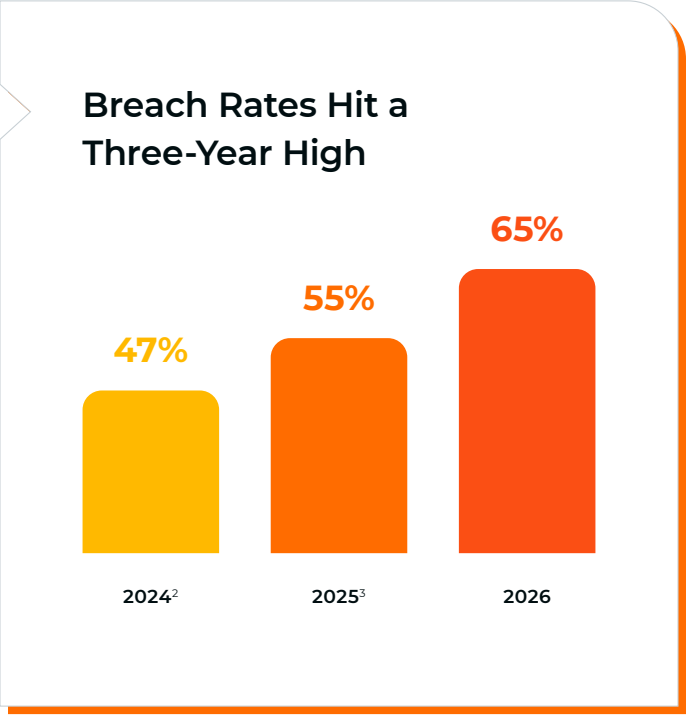

Key Findings

BREACH LEVELS HIT A NEW HIGH DRIVEN BY AI

65 percent of organizations experienced a data breach in the past 12 months with **83 percent** reporting AI involvement across external attacks and internal threats such as unsanctioned AI use (shadow AI)

TRUST ISSUES SHIFT CONFIDENCE TO DATA LAKES

71 percent are reluctant to deploy AI in public cloud environments due to concerns around intellectual property leakage, while **72 percent** believe data lakes offer greater security

AI IS THE NEWEST MEMBER OF TODAY'S CYBERSECURITY TEAM

53 percent report that alert triage and prioritization is now initiated without human interaction augmenting lean security teams

SECURITY INVESTMENTS INCREASE, BUT OUTCOMES LAG

93 percent invested in new security technologies to improve detection and visibility, yet **41 percent** report it takes longer to detect breaches

VISIBILITY EMERGES AS THE DEFINING REQUIREMENT FOR SECURITY SUCCESS

Complete visibility across all data in motion is **ranked the top** capability for improving security outcomes

While nearly **40 percent** of organizations report operating at an integrated level of AI security maturity, with practices embedded across the organization, the findings reveal a different story, with persistent gaps in governance, visibility, and skills.

That contradiction is difficult to ignore.

It reflects a growing disconnect between what organizations believe about their security posture and what they can verify.

What emerges is a false sense of confidence that hybrid cloud environments are secure. Security efficacy is being measured by what has been implemented rather than what can be proven. Tools are deployed, policies are written, and investment is increasing. But without the ability to observe how data moves, how AI behaves, and how threats develop, those indicators can be dangerously misleading.

This misplaced trust is the **AI security illusion.**

The illusion is not created by a lack of effort. It is created by a lack of visibility across increasingly complex hybrid cloud infrastructure. As environments become more distributed and AI becomes more deeply integrated, the ability to validate security is not keeping pace with the complexity of what must be secured.

At the same time, external alerts reinforce the urgency. The World Economic Forum has identified AI-driven cyber threats as a growing systemic risk[4], while global standards organizations warn that encrypted data harvested today may be decrypted in the future as quantum capabilities mature. The implications are clear. The challenge is not only current visibility, but future exposure.

This report examines where the illusion is taking hold, how it is being reinforced across hybrid cloud environments, and what it will take for organizations to move from assumption to evidence.

Because the gap between what organizations believe and what they can prove is where the AI security illusion thrives.

Confidence Without Control

Most organizations believe they are making real progress securing AI. They describe their capabilities as defined or integrated, governance frameworks are evolving, and security investments continue to grow.

Self-Reported AI Security Maturity Levels

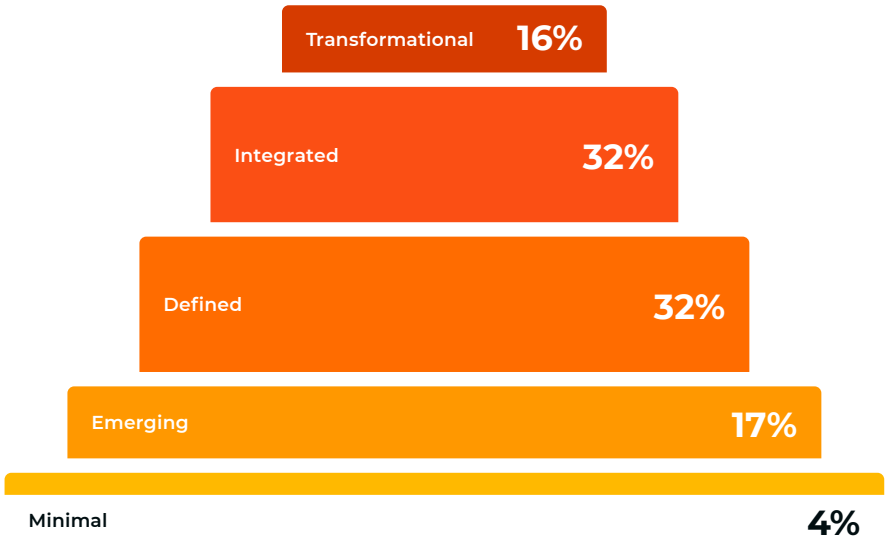

- **Transformational** Organization-wide monitoring and governance with proactive, dynamic controls
- **Integrated** Select business units adhere to uniform monitoring and robust governance
- **Defined** Established security guidelines and governance policies ensuring compliance
- **Emerging** Minimal guardrails and basic governance guidelines are forming
- **Minimal** No data monitoring and minimal governance

But these signals of maturity do not necessarily translate into control.

Organizations believe they are advancing their security programs, but with a breach rate that has grown by nearly **40 percent** in just three years[5], it's clear that not all programs are working as intended.

This challenge becomes more visible in governance.

Policies are expanding, yet **nearly half** (47 percent) of organizations report a rise in AI-related insider threats, including data leaks and unsanctioned AI use. In response, **43 percent** are now prioritizing stronger controls over employee AI use. However, as AI adoption accelerates, concerns persist. **More than three-quarters** believe unsanctioned AI use is one of the biggest barriers to secure AI adoption in the coming year.

Together, these findings reveal the challenge security teams face as they seek to leverage the innovation of AI, without compromising security. AI adoption is outpacing governance, creating exposure in areas that are difficult to monitor or control.

The same pattern appears in workforce skills.

More than three-quarters (76 percent) of Security and IT leaders report a shortage of AI security expertise on their teams, compounding an already constrained cybersecurity talent pool. At the same time, adoption of AI-driven operations is accelerating. **Nearly all** (94

percent) report using AI systems to initiate security functions without human interaction, most commonly in alert triage and prioritization (53 percent).

Organizations are responding by investing in more tools, rather than closing the skills gap.

More than **90 percent** report expanding their security stacks, and **nearly half** (48 percent) are prioritizing AI-driven tools to augment their teams. This reflects a growing reliance on automation in the absence of sufficient human expertise to guide it.

But more tools generate more data, not necessarily more clarity. Information remains distributed across systems that do not share context. Investigations become more complex, and root cause remains difficult and time consuming to determine.

What appears to be stronger coverage often creates and masks fragmentation.

AI is also changing the nature of risk. It introduces new pathways for data exposure through models, APIs, and user interactions that are harder to trace, while enabling attackers to move faster and adapt more easily.

AI-related security incidents reflect this complexity. When asked if AI was involved in the attempts or attacks organizations experienced over the past 12 months, **41 percent** cited it was an external AI-driven attack, **30 percent** stemmed from internal leaks, another **30 percent** from unsanctioned use (often referred to as shadow AI), and **one-third** targeted AI or LLM systems directly.

AI-driven attacks do not present a single risk category. Instead, organizations are recognizing that these attacks are a multi-dimensional problem both internal and external to their organizations. Without the ability to validate how systems operate, how data moves, and how threats unfold, organizations are left making assumptions about their security posture.

Over time, those assumptions harden into certainty, and that's when confidence begins to outpace control.

AI Security Incidents Span Multiple Risk Categories

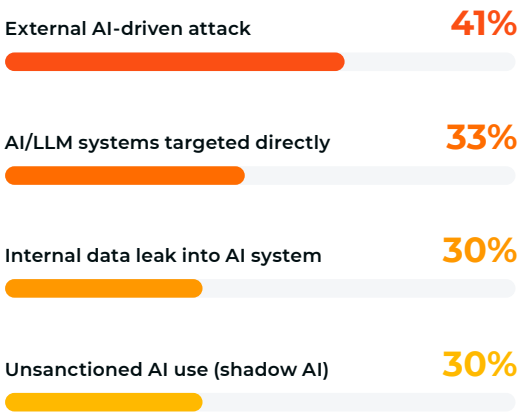

The Visibility Gap Behind the Illusion

At the center of the AI security illusion lies a critical issue: Organizations cannot see enough of what is happening inside their environments. As AI adoption accelerates, so does the volume and complexity of data in motion. Applications are no longer static. Workloads shift continuously across hybrid cloud environments. Data moves across systems, regions, and architectures in ways that are increasingly dynamic and difficult to trace. What was once a relatively contained environment is now a constantly evolving system of interactions.

That movement is where risk lives, yet most security approaches were not designed to follow it.

To keep pace with growing complexity, organizations have focused on collecting more data. Metrics, events, logs, and traces, known as MELT, continue to expand. Yet this data is not translating into greater confidence, with **nearly 4 in 10** leaders citing a lack of accurate MELT data feeding their tools as a key challenge in managing risk and accountability. The assumption is straightforward: more data should lead to better visibility. In practice, it does not.

More data does not create better understanding.

Each security tool captures a different perspective: one observes endpoints, another monitors application behavior, and another generates alerts after an event has already occurred. While each perspective is valuable, they rarely connect into a unified view of how activity actually unfolds across environments.

The result is fragmentation.

They can see that something happened, but not always how or why it happened. They can detect anomalies but not trace the full path of an interaction. They can identify symptoms, but struggle to determine cause.

As environments become more distributed across cloud, data lakes, containers, and data centers, fragmentation becomes harder to manage, and this year's research reinforces that. **Nearly half** (46 percent) of organizations report struggling with a surge in AI-powered attacks. At the same time, **45 percent** cite expanding visibility gaps driven by cloud complexity, and **43 percent** point to a shortage of cloud security expertise.

These pressures are converging in a measurable way. More than **40 percent** say it now takes longer to detect and investigate breaches, reinforcing a pattern where increased data volume without shared context is slowing, not accelerating, response.

This points to a deeper issue. The challenge is not volume. It is context.

More than ninety percent (92 percent) of organizations now agree that data security depends on complete visibility across all data in motion. Yet most organizations lack a unified view. Visibility is spread across tools, environments, and teams, making it difficult to understand how events connect or how risk develops over time.

AI amplifies this gap.

Nearly three-quarters of organizations report limited visibility into AI-driven data flows. These interactions are dynamic by design. They span models, APIs, and distributed systems, often without clear boundaries. Traditional approaches struggle to capture how data is accessed, transformed, and exposed within these workflows. Security teams recognize the need to adapt, yet **64 percent** cite insufficient budget as a major barrier, raising the risk that AI adoption will outpace their ability to secure underlying data.

THE AI BLIND SPOT

Nearly three-quarters of organizations report limited visibility into AI-driven data flows. Security teams can see that something happened, but not always how or why.

In many cases, organizations know AI is being used, but they do not fully understand how it behaves.

Encryption adds yet another layer of complexity, one that can obscure risk as easily as it protects it.

Today, malware accounts for **86 percent** of encrypted attacks, totaling a staggering 27.8 billion hits in a 12-month period[6]. Yet **three-quarters** (76 percent) of respondents still view encrypted data as inherently secure. This highlights a growing disconnect between perceived protection and reality. Encryption is widely trusted but not widely understood or monitored.

Looking ahead, this becomes more consequential in the context of quantum computing. As organizations plan for a post-quantum future, **9 in 10** (91 percent) report visibility into their encrypted traffic is critical for post quantum cryptography (PQC) readiness. At the same time, **87 percent** expressed concern about "harvest now, decrypt later" scenarios as quantum computing capabilities progress.

Global risk organizations, including the World Economic Forum[7], have reinforced this concern, warning that adversaries may already be collecting encrypted data today with the intent to decrypt it as early as 2029. In this context, visibility is not just about current detection. It is about long-term exposure.

Hybrid cloud environments amplify all of these challenges. Visibility is inconsistent across data centers, virtualized infrastructure, public cloud, and containerized environments. What is observable in one environment may be opaque in another. Data moves across these boundaries, but visibility does not move with it.

> Fragmented visibility isn't a technical inconvenience; it's a direct business liability that opens the door to breaches, compliance failures, and reputational damage. Until organizations have visibility into all data in motion, they are not managing risk, they are accepting it. In today's threat environments, that's a gamble organizations can no longer afford.

SHANE BUCKLEY
Chief Executive Officer

Top AI-Driven Threats to Hybrid Cloud Security

1. **Increase in AI-driven attacks**
2. **Visibility gaps due to cloud complexity**
3. **Shortage of cloud security expertise**
4. **Complexity driven by AI adoption**
5. **Fragmented security tools**

PRESSURES ARE CONVERGING

Leaders were asked to stack-rank the top AI-driven threats, yet the top three challenges are clustered within **3 percentage points** of each other. Organizations are not dealing with a single dominant threat. They are managing simultaneous, compounding pressures across AI attacks, visibility, and talent.

As a result, security teams are forced to reconstruct events by correlating signals across disconnected systems. They approximate behavior rather than directly observing it. They make decisions based on partial insight.

This is the structural gap behind the AI security illusion.

Closing that gap requires more than collecting data. It requires connecting it in a way that reflects how modern environments actually behave.

By combining network-derived telemetry with MELT data, security teams gain a complete view of data in motion. It allows organizations to trace how data moves across systems, understand how AI systems interact with that data, and observe how threats develop in real time. Leaders recognize this shift, with **9 out of 10** reporting that network-derived telemetry, consisting of packets, flows, and application metadata, is critical to securing data in motion today.

Instead of piecing together fragments, security teams can follow the full path of an interaction. Instead of inferring behavior, they can observe it. Instead of approximating risk, they can validate it.

As environments continue to evolve, that distinction becomes more important.

When Hybrid Cloud Risk Is Easy to Misread

When visibility is inconsistent, risk is no longer measured objectively. It is interpreted.

As AI workloads expand, organizations are reevaluating their data strategies to determine where data should reside across their environments. This year's findings point to an accelerating unease with the public cloud. Nearly **8 in 10** (78 percent) now view public cloud as riskier for AI, an increase of **11 percentage points** year over year. More than **70 percent** express reluctance to deploy AI in public cloud environments citing specific concerns around intellectual property leakage, up sharply from **54 percent** the previous year.

In response, many organizations are shifting their assumption about where data is safest. **Nearly three-quarters** (72 percent) now believe data lakes are inherently more secure for their AI workloads.

These shifts are significant. They reflect a reassessment of where organizations believe they can best maintain visibility and control over AI data.

That shift is translating into action. **Nearly 80 percent** are considering repatriating workloads from public cloud environments, up from **70 percent** in the prior year. At the same time, the urgency to secure cloud infrastructure persists, with Gartner forecasting a jump in cloud security spending from $8.98B in 2024 to $26.2B in 2029[8].

Yet even as strategies evolve, the underlying challenge remains unchanged. Risk is not determined by where data resides, but by whether it can be consistently observed, understood, and secured across environments.

The AI Security Illusion Reaches the Boardroom

The gap between perception and verification does not remain confined to security operations. It extends into how organizations make decisions at the highest levels.

At the board level, security is often evaluated through indicators that suggest progress: rising investment, expanding toolsets, and more formalized governance frameworks. Together, these signals create a sense that risk is being addressed, and environments are becoming more secure.

But those indicators do not always reflect operational reality.

More than two-thirds (67 percent) of organizations cite a lack of board-level understanding of AI-related security risks as a major barrier, increasing the risk that AI adoption will outpace security. This is not a lack of engagement. In many cases, boards are more involved than ever. The issue is alignment. Security is being interpreted through a strategic lens, while the underlying challenges remain deeply operational and often invisible.

That disconnect becomes most visible during incidents.

The 72-hour Illusion

Nearly half of CIOs and CTOs believe root cause can be identified within 72 hours. Only **27 percent** of CISOs agree, with **nearly half** saying it takes up to 7 days, and **almost a quarter** reporting it can take up to 30 days. This misalignment leaves CISOs under pressure to act before incidents are fully understood.

Nearly half of C-level executives surveyed (CTOs and CIOs) believe root cause of security incidents can be identified within 72 hours. In contrast, only about **one-quarter** of CISOs agree, with **nearly half** reporting that identification takes up to seven days. In practice, investigations frequently take longer, particularly in environments where AI systems, distributed architectures, and fragmented visibility complicate analysis.

The result is a misalignment between boardroom expectations and operational reality, leaving CISOs under pressure to act on incidents before they are fully understood.

That pressure is grounded in operational reality. CISOs report that identification of root cause of a breach can take several days, highlighting the complexity of investigating incidents in modern, distributed environments. Determining how data moved, how systems interacted, and where exposure occurred requires a level of visibility that many organizations do not yet have.

Without that visibility, investigations become slower, less certain, and more resource-intensive. The consequences are tangible.

More than 40 percent of organizations report financial losses tied to security incidents. **Over one-third** (37 percent) cite increased cyber insurance premiums following breaches, and **nearly one-third** (32 percent) report regulatory or compliance penalties. These are not isolated outcomes, but indicators of a broader shift in how cyber risk translates into business impact. At the same time, regulatory expectations are increasing.

Frameworks such as NIS2 and DORA are raising the bar for operational resilience and accountability. GDPR enforcement continues to drive significant financial penalties for failures in data protection. In the United States, the SEC is placing greater emphasis on transparency and timely disclosure of cyber incidents. Across regions, organizations are being held to higher standards for how they manage and report risk.

This raises the stakes for alignment.

THE COST OF THE ILLUSION

In addition to direct financial costs of breaches, including increased cyber insurance premiums and regulatory penalties, **34 percent** of Security and IT leaders reported loss of corporate or customer data from their most serious breach. This loss has far-reaching financial, operational, and reputational consequences, further elevating cybersecurity to a board-level priority.

CISO PERSPECTIVE

From Confidence to Proof

The AI security illusion is most visible at the leadership level, where perception and operational reality begin to diverge.

For CISOs, the stakes are both organizational and personal. **More than 1 in 4** are concerned about losing their job as a result of a cyber incident, underscoring how the gap between perceived control and actual capability translates directly into personal accountability.

This pressure is compounded by limited visibility. **Nearly 9 in 10** CISOs (89 percent) believe that improved visibility into encrypted traffic would reduce cyber insurance premiums, and **85 percent** say application metadata is critical to securing encrypted traffic. Without deeper, more reliable insight into how environments behave, CISOs are held responsible for risks they cannot fully observe.

The issue is not simply scale. It is evidence.

Breaking the illusion requires a shift from monitoring activity to understanding behavior and from implementing controls to proving outcomes. Access to complete, accurate insight across data in motion is essential to closing the gap between confidence and control.

At the same time, CISOs must translate that reality into terms the business can act on. **Seven in 10** CISOs believe the board's lack of understanding of security best practices could cause AI adoption to outpace security readiness.

This is reshaping priorities. **Forty-one percent** of CISOs cite improving board-level understanding of the risks and benefits of AI as a top security priority over the next 12 months, while **more than a quarter** identify improved reporting aligned to business outcomes as the most important step in strengthening board engagement.

What CISOs Need Most from the Board

- **Improving board-level reporting by demonstrating security alignment with business outcomes**
- **Ensuring cybersecurity is a critical element on the board's risk agenda**
- **Support for CISOs on responsibility and accountability for security posture**

Forrester reinforces this shift, highlighting alignment between security leadership and business strategy as a defining priority for modern CISOs[9].

Without that alignment, the illusion persists.

Boards see investment and assume progress. Security teams see fragmentation and uncertainty. Both perspectives are valid, but they are not reconciled. Over time, this disconnect shapes decision making, and resources are allocated based on perceived coverage rather than verified capability. Risk is evaluated based on high-level indicators rather than operational evidence. AI adoption continues to accelerate, while the ability to validate security lags.

Closing that gap requires a shared foundation of visibility and evidence. When organizations can observe how their environments behave in real time, they can align technical reality with strategic decision-making.

Until then, the AI security illusion remains both a technical and leadership challenge.

This perspective is explored further in the **2026 CISO Perspectives report**.

Where Security Leaders are Focusing Next

VISIBILITY

Improving visibility into AI-driven data flows and traffic across hybrid cloud environments

TOOLING

Using AI-powered security tools to augment internal security teams and workflows

OPERATIONS

Strengthening incident detection and response capabilities for AI-related threats

GUARDRAILS

Implementing security guardrails and monitoring controls around AI and LLM usage

GOVERNANCE

Developing stronger corporate governance for the appropriate use of AI by employees

From Illusion to Evidence

The survey does not just expose the gap between perception and reality. It also points to how organizations can begin to close it.

At its core, the challenge is not a lack of intent or effort. Organizations are actively evolving their security strategies, expanding governance, and adopting new technologies. **All 1,023 respondents** report that AI security tools are having a positive impact within their organizations, particularly in areas such as automation, productivity, and threat analysis. However, these tools are not yet grounded in a complete understanding of how their environments behave. Security is being inferred rather than observed.

To move forward, organizations need to shift from fragmented visibility to connected insight. They need the ability to understand how data moves across hybrid cloud environments, how AI systems interact with that data, and how threats develop in real time.

This is the foundation for moving from assumption to evidence.

Deep observability changes the model.

More than 90 percent of organizations say deep observability is foundational to securing AI deployments. At the same time, as confidence in data lakes grows, **92 percent** report that network-derived telemetry is critical to securing data in motion. This reflects a broader recognition that traditional visibility approaches are no longer sufficient for the scale and complexity of modern environments.

By combining network-derived telemetry with MELT data, organizations gain a unified view of data in motion. Instead of analyzing isolated signals, they can follow the full path of an interaction across systems, environments, and workflows.

This shift has both practical and positive implications.

Security teams can validate how controls are performing rather than assuming they are effective. They can trace the origin and progression of an incident

rather than reconstructing it from fragments. They can identify exposure points across AI systems, APIs, and distributed architectures with greater precision.

Security leaders are clear about what is needed to enable this shift.

They point to accurate network-derived telemetry, comprehensive visibility into all data in motion, including encrypted traffic, and the resources required to interpret and act on that data. These are not incremental improvements, they are foundational capabilities.

Deep observability does not replace existing tools. It connects them and provides the context needed to understand how systems interact, how data flows, and where risk emerges. It turns fragmented visibility into a coherent, actionable view.

Ultimately, the goal is not simply better detection or faster response, it is confidence grounded in evidence-based proof.

Because in an environment defined by AI, hybrid cloud, and accelerating complexity, the organizations that can validate their security posture will be the ones that can trust it.

And that is how AI security moves from illusion to reality.

Methodology

The data presented in this report was compiled by Vitreous World. Fieldwork was conducted using an online methodology, recruiting a mix of Chief Information Officers, Chief Information Security Officers, Chief Technology Officers, Chief Risk Officers and those working in information technology, cybersecurity or security operations, information security and other technology roles were recruited. Interviews were conducted across Australia, France, Germany, Singapore, the UK and the USA.

All respondents were guaranteed to remain anonymous as part of the study. Fieldwork was carried out between February 16-17, 2026.

1023 Respondents

46% Work for companies with **between 501 and 1,000 employees**

54% Work for companies with **more than 1,000 employees**

61% **Senior management** (C-Suite, C-Level)

39% **Middle management** (Vice President, Director, Department Head, Senior Manager)

622 Senior Leaders

49% CISO

31% Chief Technology

20% Chief Information Officer

Countries

160 Australia

201 France

150 Germany

160 Singapore

151 UK

201 USA

About Gigamon

Gigamon® protects the hybrid cloud networks and data of the world's most complex organizations. The AI-powered Gigamon Deep Observability Pipeline delivers complete visibility into all data in motion by providing trusted, network-derived telemetry directly to cloud, security, and observability tools. With AI-driven insights across packets, flows, and application metadata, organizations can detect threats concealed in encrypted and lateral traffic, resolve network and application performance bottlenecks, and validate compliance while reducing cost and complexity. Gigamon is trusted by over 4,000 organizations worldwide, including 83 of the Fortune 100, major mobile network operators, and public sector agencies at every level.

Learn more at **gigamon.com**.

References

1 Gartner, 2025, Gartner Says Worldwide AI Spending Will Total $1.5 Trillion in 2025

2 Gigamon, 2024, Hybrid Cloud Security Survey: Closing the Cybersecurity Preparedness Gap

3 Gigamon, 2025, Hybrid Cloud Security Survey: Evolving Hybrid Cloud Security in the Age of AI

4 World Economic Forum, 2026, Cyber risk in 2026: What executives must know about AI, fraud, geopolitics and more

5 Gigamon, 2024, Hybrid Cloud Security: Closing the Cybersecurity Preparedness Gap

6 Zscaler, 2024, Zscaler Finds Over 87% of Cyberthreats Hide in Encrypted Traffic, Reinforcing Need For Zero Trust

7 World Economic Forum, 2026, Why quantum security is a question leaders cannot ignore right now

8 Gartner, 2025, Information Security Forecast.

9 Forrester, 2025, High-Performance IT Makes Business Results Technology's North Star

Gigamon®

Worldwide Headquarters
3300 Olcott Street, Santa Clara, CA 95054 USA
+1 (408) 831-4000 | gigamon.com

04.26_01